Approximate Maximum A Posteriori Inference with Entropic Priors

Matthew D. Hoffman
Department of Statistics
Columbia University
New York, NY
mdhoffma@cs.princeton.edu

Abstract

In certain applications it is useful to fit multinomial distributions to observed data with a penalty term that encourages sparsity. For example, in probabilistic latent audio source decomposition one may wish to encode the assumption that only a few latent sources are active at any given time. The standard heuristic of applying an L1 penalty is not an option when fitting the parameters to a multinomial distribution, which are constrained to sum to 1. An alternative is to use a penalty term that encourages low-entropy solutions, which corresponds to maximum a posteriori (MAP) parameter estimation with an entropic prior. The lack of conjugacy between the entropic prior and the multinomial distribution complicates this approach. In this report I propose a simple iterative algorithm for MAP estimation of multinomial distributions with sparsity-inducing entropic priors.

1 Introduction

Suppose we want to estimate the parameter $\boldsymbol{\theta}$ to a multinomial distribution responsible for generating N observations $x_i \in \{1, \ldots, K\}$. The log-likelihood of the data is given by

$$\log p(\boldsymbol{x}) = \sum_i \log \theta_{x_i}, \tag{1}$$

and the maximum-likelihood estimate of $\boldsymbol{\theta}$ is simply

$$\theta_k \propto \sum_i \mathbb{I}[x_i = k], \tag{2}$$

where $\mathbb{I}$ is an indicator function whose value is 1 if its argument is true and 0 if its argument is false.

The maximum-likelihood estimate may not be optimal if we have a priori knowledge that leads us to believe that $\boldsymbol{\theta}$ is sparse. For example, if $\boldsymbol{\theta}$ indicated the relative loudness of a set of 88 piano notes at particular moment (as it might in an application of Probabilistic Latent Component Analysis to audio spectrograms [1]), then we might expect only a few elements of $\boldsymbol{\theta}$ to be much greater than 0. This would correspond to the intuition that pianists rarely mash the entire piano keyboard at once.

To incorporate this prior intuition into our analysis, we might add a penalty term to our log-likelihood function that encourages sparse settings of $\boldsymbol{\theta}$. A common heuristic for inducing sparsity in optimization problems is to introduce an L1 penalty term into the cost function (for example, in lasso regression [2]). This is not an option here, since the L1 norm of $\boldsymbol{\theta}$ is constrained to be 1. A natural alternative is to include a negative entropy term in the log-likelihood function, corresponding to placing an unnormalized sparse entropic prior on $\boldsymbol{\theta}$:

$$\log p(\boldsymbol{\theta}, \boldsymbol{x}) = \text{constant} + a \sum_k \theta_k \log \theta_k + \sum_i \log \theta_{x_i}. \tag{3}$$

The constant a controls the strength of the prior $p(\boldsymbol{\theta}) \propto \exp\{a \sum_k \theta_k \log \theta_k\}$. If a is positive, then this prior will give higher weight to low-entropy settings of $\boldsymbol{\theta}$.

Unfortunately, the Maximum A Posteriori (MAP) estimate of $\boldsymbol{\theta}$ does not have a simple analytic form for this model, since the entropic prior is not conjugate to the multinomial distribution. In the following section, I propose a simple iterative scheme for MAP estimation of $\boldsymbol{\theta}$ when a is positive.

2 A MAP Inference Scheme for the Sparse Entropic Prior

Our strategy is based on optimizing the following approximate auxiliary function for the negative entropy term in equation 3:

$$\ell(a, \nu, \boldsymbol{\theta}, \boldsymbol{\alpha}) \triangleq a \sum_k \alpha_k (\nu \log \theta_k - (\nu - 1) \log \alpha_k), \tag{4}$$

where $\boldsymbol{\alpha}$ is a free parameter such that $\sum_k \alpha_k = 1$ and $\alpha_k \geq 0$, and ν is a real-valued scalar constrained to be greater than 1. Taking the derivative of the Lagrangian of ℓ with respect to α_k yields

$$\frac{\partial \ell}{\partial \alpha_k} = a\nu \log \theta_k - a(\nu - 1)(1 + \log \alpha_k) + \lambda. \tag{5}$$

Setting the right side equal to zero shows that ℓ is optimized with respect to $\boldsymbol{\alpha}$ when

$$\alpha_k \propto \exp\left\{\frac{\nu}{\nu - 1} \log \theta_k\right\} = \theta_k^{\frac{\nu}{\nu-1}}. \tag{6}$$

When ν is large, this implies that the optimal value of ℓ can only be achieved when $\alpha_k \approx \theta_k$.

When $\alpha_k = \theta_k$, we recover the original entropic prior term:

$$\ell(a, \nu, \boldsymbol{\theta}, \boldsymbol{\theta}) = a \sum_k \theta_k (\nu \log \theta_k - (\nu - 1) \log \theta_k) = a \sum_k \theta_k \log \theta_k. \tag{7}$$

Thus, for sufficiently large values of ν, when $\boldsymbol{\alpha}$ is optimally chosen ℓ approximates the entropic prior. We may therefore substitute ℓ (with a large value of ν) for the entropic prior term in equation 3 and jointly optimize the approximate objective

$$\mathcal{L} \triangleq a \sum_k \alpha_k (\nu \log \theta_k - (\nu - 1) \log \alpha_k) + \sum_i \log \theta_{x_i}. \tag{8}$$

over $\boldsymbol{\alpha}$ and $\boldsymbol{\theta}$. When the gradient of $\mathcal{L}$ with respect to $\boldsymbol{\alpha}$ is 0, as it must be at a local optimum of $\mathcal{L}$, $\alpha_k \approx \theta_k$, and so $\mathcal{L} \approx \log p(\boldsymbol{x}, \boldsymbol{\theta})$, the objective function of interest.

A simple fixed-point iteration can be used to optimize $\mathcal{L}$ over $\boldsymbol{\alpha}$ and $\boldsymbol{\theta}$. The gradient of the Lagrangian with respect to θ_k is

$$\frac{\partial \mathcal{L}}{\partial \theta_k} = \frac{1}{\theta_k}\left(a\alpha_k\nu + \sum_i \mathbb{I}[x_i = k]\right) + \lambda, \tag{9}$$

where $\mathbb{I}$ is an indicator function whose value is 1 if its argument is true and 0 if its argument is false. $\mathcal{L}$ is therefore maximized with respect to $\boldsymbol{\theta}$ when

$$\theta_k \propto a\alpha_k\nu + \sum_i \mathbb{I}[x_i = k]. \tag{10}$$

As observed above, $\mathcal{L}$ is maximized with respect to $\boldsymbol{\alpha}$ when

$$\alpha_k \propto \theta_k^{\frac{\nu}{\nu-1}}. \tag{11}$$

By iterating between the updates in equation 10 and 11, we reach a stationary point of $\mathcal{L}$. At such a stationary point, $\mathcal{L} \approx \log p(\boldsymbol{x}, \boldsymbol{\theta})$, and so we may conclude that the value of $\boldsymbol{\theta}$ at a stationary point of $\mathcal{L}$ yields approximately a local optimum of $\log p(\boldsymbol{x}, \boldsymbol{\theta})$.

Note that these updates may require a number of iterations to converge, and the number of iterations needed is likely to grow with ν. However, the cost of each update is minimal. If these updates are incorporated as part of a larger coordinate ascent algorithm like the expectation-maximization algorithm used in probabilistic latent semantic indexing [3], the additional expense involved in iterating between updating $\boldsymbol{\theta}$ and $\boldsymbol{\alpha}$ is likely to be dominated by the cost of computing $\sum_i \mathbb{I}[x_i = k]$ (or its expected value).

3 Conclusion

I have presented a simple fixed-point iteration for performing approximate maximum a posteriori estimation of multinomial parameters in the presence of a sparsity-inducing entropic prior. This algorithm only provides an approximate solution, but it can be made arbitrarily accurate at the cost of slower convergence. The algorithm is very easy to implement, and the cost per iteration is minimal.

References

[1] P. Smaragdis, B. Raj, and M. Shashanka. A probabilistic latent variable model for acoustic modeling. In *Workshop on Advances in Models for Acoustic Processing at NIPS*, 2006.

[2] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning.* Springer, August 2001.

[3] T. Hofmann. Probabilistic latent semantic analysis. In *Uncertainty in Artificial Intelligence (UAI)*, 1999.